arXiv:1409.2979v4 [math.OC] 29 Oct 2014

Proximal Stochastic Newton-type Gradient Descent Methods for Minimizing Regularized Finite Sums

Ziqiang Shi*†

August, 2014

Abstract

In this work, we generalized and unified recent two completely different works of Jascha [9] and Lee [2] respectively into one by proposing the **prox**imal s**to**chastic **N**ewton-type gradient (PROXTONE) method for optimizing the sums of two convex functions: one is the average of a huge number of smooth convex functions, and the other is a non-smooth convex function. While a set of recently proposed proximal stochastic gradient methods, include MISO, Prox-SDCA, Prox-SVRG, and SAG, converge at linear rates, the PROXTONE incorporates second order information to obtain stronger convergence results, that it achieves a linear convergence rate not only in the value of the objective function, but also in the *solution*. The proof is simple and intuitive, and the results and technique can be served as a initiate for the research on the proximal stochastic methods that employ second order information.

1 Introduction and problem statement

In this work, we consider the problems of the following form:

$$\underset{x\in\mathbf{R}^p}{\text{minimize}}\ f(x) := \frac{1}{n}\sum_{i=1}^{n} g_i(x) + h(x), \tag{1.1}$$

where g_i is a smooth convex loss function associated with a sample in a training set, and h is a non-smooth convex penalty function or regularizer. Let $g(x) = \frac{1}{n}\sum_{i=0}^{n} g_i(x)$. We assume the optimal value $f^\star$ is attained at some optimal solution $x^\star$, not necessarily unique. Problems of this form often arise in machine learning, such as the least-squares regression, the Lasso, the elastic net, and the logistic regression.

For optimizing (1.1), the standard and popular *proximal full gradient method* (Prox-FG) uses iterations of the form

$$x^{k+1} = \underset{x\in\mathbf{R}^p}{\arg\min}\left\{\nabla g(x_k)^T x + \frac{1}{2\alpha_k}\|x - x_{k-1}\|^2 + h(x)\right\}, \tag{1.2}$$

where α_k is the step size at the k-th iteration. Under standard assumptions the sub-optimality achieved on iteration k of the Prox-FG method with a constant step size is given by

$$\mathbb{E}[f(x^k)] - f(x^*) = O(\frac{1}{k}).$$

When f is strongly-convex, the error satisfies [10]

$$\mathbb{E}[f(x^k)] - f(x^*) = O((\frac{L-\mu_g}{L+\mu_h})^k),$$

where L is the Lipschitz constant of $f(x)$, μ_g, and μ_h are the convexity parameters of $g(x)$ and $h(x)$ respectively. These notations will be detailed in Section 1.1. This results in a linear convergence rate, which is also known as a geometric or exponential rate because the error is cut by a fixed fraction on each iteration.

*Fujitsu Research & Development Center, Beijing, China.
†shiziqiang@cn.fujitsu.com; shiziqiang7@gmail.com.

Unfortunately, the Prox-FG and methods can be unappealing when n is large because its iteration cost scales linearly in n. When the number of components n is very large, each iteration of (1.2) can be very expensive since it requires computing the gradients for all the n component functions g_i, and also their average.

The main appeal of proximal stochastic gradient (Prox-SG) methods is that they have an iteration cost which is independent of n, making them suited for modern problems where n may be very large. The basic Prox-SG method for optimizing (1.1), uses iterations of the form

$$x_k = \text{prox}_{\alpha_k h}\big(x_{k-1} - \alpha_k \nabla g_{i_k}(x_{k-1})\big), \tag{1.3}$$

where at each iteration an index i_k is sampled uniformly from the set $\{1, ..., n\}$. The randomly chosen gradient $\nabla g_{i_k}(x_{k-1})$ yields an unbiased estimate of the true gradient $\nabla g(x_{k-1})$ and one can show under standard assumptions that, for a suitably chosen decreasing step-size sequence $\{\alpha_k\}$, the Prox-SG iterations have an expected sub-optimality for convex objectives of [1]

$$\mathbb{E}[f(x^k)] - f(x^*) = O(\frac{1}{\sqrt{k}})$$

and an expected sub-optimality for strongly-convex objectives of

$$\mathbb{E}[f(x^k)] - f(x^*) = O(\frac{1}{k}).$$

In these rates, the expectations are taken with respect to the selection of the i_k variables.

There is another group of methods, which converges much faster, but need more memory and computation to obtain the second order information about the objective function. These methods are always limited to small-to-medium scale problems that require a high degree of precision. For optimizing (1.1), *proximal Newton-type* methods [2] that incorporate second order information use iterations of the form $x^{k+1} \leftarrow x^k + \Delta x^k$, here Δx^k is obtained by

$$\Delta x^k = \underset{d \in \mathbf{R}^p}{\arg\min}\ \nabla g(x^k)^T d + \frac{1}{2} d^T H_k d + h(x^k + d), \tag{1.4}$$

where H_k denotes an approximation to $\nabla^2 g(x_k)$. According to the strategies for choosing H_k, we obtain different method, such as *proximal Newton method* (Prox-N) when we choose H_k to be $\nabla^2 g(x^k)$; *proximal quasi-Newton method* (Prox-QN) when we build an approximation to $\nabla^2 g(x_k)$ using changes measured in ∇g according to a quasi-Newton strategy [2]. Indeed if we compared (1.4) with (1.2), it can be seen Prox-N is the Prox-FG with scaled proximal mappings.

Based on the background above, now we can describe our approaches and findings. The primary contribution of this work is the proposal and analysis of a new algorithm that we call the proximal stochastic Newton-type gradient (PROXTONE, pronounced /prok stone/) method, a stochastic variant of the Prox-N method. The PROXTONE method has the low iteration cost as that of Prox-SG methods, but achieves the convergence rates stated above for the Prox-FG method. The PROXTONE iterations take the form $x^{k+1} \leftarrow x^k + t_k \Delta x^k$, where Δx^k is obtained by

$$\Delta x^k \leftarrow \arg\min_d d^T(\nabla_k + H_k x^k) + \frac{1}{2} d^T H_k d + h(x^k + d), \tag{1.5}$$

here $\nabla_k = \frac{1}{n}\sum_{i=1}^n \nabla_k^i$, $H_k = \frac{1}{n}\sum_{i=1}^n H_k^i$, and at each iteration a random index j and corresponding H_{k+1}^j is selected, then we set

$$\nabla_{k+1}^i = \begin{cases} \nabla g_i(x^{k+1}) - H_{k+1}^i x^{k+1} & \text{if } i = j, \\ \nabla_{k+1}^i & \text{otherwise.} \end{cases}$$

and $H_{k+1}^i \leftarrow H_k^i$ $(i \neq j)$.

That is, like the Prox-FG and Prox-N method, the step incorporates a gradient with respect to each function. But, like the Prox-SG method, each iteration only computes the gradient with respect to a single example and the cost of the iterations is independent of n. Despite the low cost of the PROXTONE iterations, we show in this paper that the PROXTONE iterations have a linear convergence rate for strongly-convex

objectives, like the Prox-FG method. That is, by having access to j and by keeping a memory of the approximation for the Hessian matrix computed for the objective funtion, this iteration achieves a faster convergence rate than is possible for standard Prox-SG methods.

There are a large variety of approaches available to accelerate the convergence of Prox-SG methods, and a full review of this immense literature would be outside the scope of this work. Several recent work considered various special cases of (1.1), and developed algorithms that enjoy the linear convergence rate, such as Prox-SDCA [7], MISO [3], SAG [6], Prox-SVRG [10], SFO [9], and Prox-N [2]. All these methods converge with an exponential rate in the value of the objective function, except that the Prox-N achieves superlinear rates of convergence in *solution*, however it is a batch mode method. Shalev-Shwartz and Zhang [8, 7]'s Prox-SDCA considered the case where the component functions have the form $g_i(x) = \phi_i(a_i^T x)$ and the Fenchel conjugate functions of ϕ_i and h can be computed efficiently. Schimidt et al. [6]'s SAG and Jascha et al. [9]'s SFO considered the case where $h(x) \equiv 0$.

Our PROXTONE is a extension of the SFO and Prox-N to a proximal stochastic Newton-type method for solving the more *general* (compared to Prox-SDCA, SAG and SFO) class of problems defined in (1.1). PROXTONE makes connections between two completely different approaches. It achieves a linear convergence rate not only in the value of the objective function, but also in the *solution*. We now outline the rest of the study. Section 2 presents the main algorithm and gives a equivalent form in order for the ease of analysis. Section 3 states the assumptions underlying our analysis and gives the main results; we first give a linear convergence rate in *function value* (weak convergence) that applies for any problem, and then give a strong linear convergence rate in *solution*, however with some additional conditions. Finally we conclude in Section 4.

1.1 Notations and Assumptions

Before proceeding, we introduce the notations and some useful lemmas formally first. In this work, we most adopt the nomenclature used by Nesterov [5]. The functions encountered in this work are all convex if there are no other statements.

In this paper, we assume the function $h(x)$ is lower semi-continuous and convex, and its effective domain, $\mathrm{dom}(h) := \{x \in \mathbf{R}^p \,|\, h(x) < +\infty\}$, is closed. Each $g_i(x)$, for $i = 1, \ldots, n$, is differentiable on an open set that contains $\mathrm{dom}(h)$, and their gradients are Lipschitz continuous. That is, there exist $L_i > 0$ such that for all $x, y \in \mathrm{dom}(h)$,

$$\|\nabla g_i(x) - \nabla g_i(y)\| \leq L_i \|x - y\|. \tag{1.6}$$

This is a fairly weak assumption on the g_i functions, and in cases where the g_i are twice-differentiable it is equivalent to saying that the eigenvalues of the Hessians of each g_i are bounded above by L_i. Then from the Lemma 1.2.3 and its proof in Nesterov's book, for $i = 1, \ldots, n$, we have

$$|g_i(x) - g_i(y) - \nabla g_i(y)^T (x - y)| \leq \frac{L_i}{2} \|x - y\|^2. \tag{1.7}$$

The above assumption of (1.6) implies that the gradient of the average function $g(x)$ is also Lipschitz continuous, i.e., there is an $L > 0$ such that for all $x, y \in \mathrm{dom}(h)$,

$$\|\nabla g(x) - \nabla g(y)\| \leq L \|x - y\|.$$

Moreover, we have $L \leq (1/n) \sum_{i=1}^n L_i$.

A function $f(x)$ is called μ-strongly convex, if there exist $\mu \geq 0$ such that for all $x \in \mathrm{dom}(f)$ and $y \in \mathbf{R}^p$,

$$f(y) \geq f(x) + \xi^T (y - x) + \frac{\mu}{2} \|y - x\|^2, \quad \forall\, \xi \in \partial f(x). \tag{1.8}$$

The *convexity parameter* of a function is the largest μ such that the above condition holds. If $\mu = 0$, it is identical to the definition of a convex function. The strong convexity of $f(x)$ in (1.1) may come from either $g(x)$ or $h(x)$ or both. More precisely, let $g(x)$ and $h(x)$ have convexity parameters μ_g and μ_h respectively, then $\mu \geq \mu_g + \mu_h$. From Lemma B.5 in [3] and (1.8), we have

$$f(y) \geq f(x^*) + \frac{\mu}{2} \|y - x^*\|^2. \tag{1.9}$$

2 The PROXTONE method

We summarize the PROXTONE method of (1.5) in Algorithm 1. It can be easily checked that if $n = 1$, then it becomes the determined proximal Newton-type methods proposed by Lee and Sun et al. [2] for minimizing composite functions:

$$\underset{x\in\mathbf{R}^p}{\text{minimize}}\ f(x) := g(x) + h(x) \tag{2.1}$$

by (1.4), thus PROXTONE is indeed a generalization of Prox-NG.

Algorithm 1 PROXTONE: A generic PROXimal sTOchastic NEwton-type gradient descent method

Input: start point $x^0 \in \text{dom } f$; for $i \in \{1, 2, .., n\}$, let $H^i_{-1} = H^i_0$ be a positive definite approximation to the Hessian of $g_i(x)$ at x^0, $\nabla^i_{-1} = \nabla^i_0 = \nabla g_i(x^0) - H^i_0 x^0$; and $\nabla_0 = \frac{1}{n}\sum_{i=1}^n \nabla^i_0$, $H_0 = \frac{1}{n}\sum_{i=0}^n H^i_0$.
1: **repeat**
2: Solve the subproblem for a search direction:$\triangle x^k \leftarrow \arg\min_d d^T(\nabla_k + H_k x^k) + \frac{1}{2}d^T H_k d + h(x^k + d)$.
3: Update: $x^{k+1} = x^k + \triangle x^k$.
4: Sample j from $\{1, 2, .., n\}$, use the $\nabla g_j(x^{k+1})$ and H^j_{k+1}, which is a positive definite approximation to the Hessian of $g_j(x)$ at x^{k+1}, to update the ∇^i_{k+1} $(i \in \{1, 2, .., n\})$: $\nabla^j_{k+1} \leftarrow \nabla g_j(x^{k+1}) - H^j_{k+1}x^{k+1}$, while leaving all other ∇^i_{k+1} and H^i_{k+1} unchanged: $\nabla^i_{k+1} \leftarrow \nabla^i_k$ and $H^i_{k+1} \leftarrow H^i_k$ $(i \neq j)$; and finally obtain ∇_{k+1} and H_{k+1} by $\nabla_{k+1} \leftarrow \frac{1}{n}\sum_{i=1}^n \nabla^i_{k+1}$, $H_{k+1} \leftarrow \frac{1}{n}\sum_{i=1}^n H^i_{k+1}$.
5: **until** stopping conditions are satisfied.
Output: x^k.

It is also a generalization of recent work by Jascha [9], whose SFO is the special case of our PROXTONE with $h(x) \equiv 0$. Our algorithm in Jascha's style is summarized in Algorithm 2 which is equivalent to the original PROXTONE. To see the equivalence, keep in mind that $G^k(x)$ is a quadratic function, we only need to check the following equations:

$$\nabla^2 G^k(x) = \frac{1}{n}\sum_{i=1}^n H^i_k \quad \text{and} \quad \nabla G^k(x) = \frac{1}{n}\sum_{i=1}^n \nabla g_i(x) + \frac{1}{n}\sum_{i=1}^n (x - x^k)^T H^i_k,$$

and

$$\nabla_k + H_k x^k = \frac{1}{n}\sum_{i=1}^n [\nabla g_i(x^{\theta_{i,k}}) + (x^k - x^{\theta_{i,k-1}})^T H^i_{\theta_{i,k}}].$$

In following analysis of Section 3, we will not distinguish these two forms from each other.

Algorithm 2 PROXTONE in a form that is easy to analyze

Input: start point $x^0 \in \text{dom } f$; for $i \in \{1, 2, .., n\}$, let $g^0_i(x) = g_i(x^0) + (x - x^0)^T \nabla g_i(x^0) + \frac{1}{2}(x - x^0)^T H^i_0 (x - x^0)$, where the notation H^i_0 $(i \in \{1, 2, .., n\})$ are totally the same as they in Algorithm 1; and $G^0(x) = \frac{1}{n}\sum_{i=1}^n g^0_i(x)$.
1: **repeat**
2: Solve the subproblem for new approximation of the solution: $x^{k+1} \leftarrow \arg\min_x \big[G^k(x) + h(x)\big]$.
3: Sample j from $\{1, 2, .., n\}$, and update the surrogate functions:

$$g^{k+1}_j(x) = g_j(x^{k+1}) + (x - x^{k+1})^T \nabla g_j(x^{k+1}) + \frac{1}{2}(x - x^{k+1})^T H^i_{k+1}(x - x^{k+1}), \tag{2.2}$$

while leaving all other $g^{k+1}_i(x)$ unchanged: $g^{k+1}_i(x) \leftarrow g^k_i(x)$ $(i \neq j)$; and $G^{k+1}(x) = \frac{1}{n}\sum_{i=1}^n g^{k+1}_i(x)$.
4: **until** stopping conditions are satisfied.
Output: x^k.

3 Convergence Analysis

Under the standard assumptions, we now state our convergence result.

Theorem 3.1. *Suppose $\nabla g_i(x)$ is Lipschitz continuous with constant $L_i > 0$ for $i = 1, ..., n$, and $L_i I \preceq mI \preceq H_k^i \preceq MI$ for all $i = 1, ..., n$ and $k \geq 1$, $h(x)$ is strongly convex with $\mu_h \geq 0$, then the PROXTONE iterations satisfy for $k \geq 1$:*

$$\mathbb{E}[f(x^k)] - f^* \leq \frac{M + L_{max}}{2}[\frac{1}{n}\frac{M + L_{max}}{2\mu_h + m} + (1 - \frac{1}{n})]^k \|x^* - x^0\|^2. \tag{3.1}$$

The ideas of the proof is near identical to that of MISO by Mairal [3] and for completeness we give a simple version in the appendix.

We have the following remarks regarding the above result:

- In order to satisfy $\mathbb{E}[f(x^k)] - f^* \leq \epsilon$, the number of iterations k needs to satisfy

$$k \geq (\log \rho)^{-1} \log[\frac{2\epsilon}{(M + L_{max})\|x^* - x^0\|^2}],$$

  where $\rho = \frac{1}{n}\frac{M+L_{max}}{2\mu_h+m} + (1 - \frac{1}{n})$.

- Inequality (3.1) gives us a reliable stopping criterion for the PROXTONE method.

At this moment, we see that the expected quality of the output of PROXTONE is good. However, in practice we are not going to run this method many times on the same problem. What is the probability that our single run can give us also a good result. Since $f(x^k) - f^* \geq 0$, Markov's inequality and Theorem 3.1 imply that for any $\epsilon > 0$,

$$\mathbf{Pr}\Big(f(x^k) - f^* \geq \epsilon\Big) \;\leq\; \frac{\mathbb{E}[f(x^k) - f^*]}{\epsilon} \;\leq\; \frac{(M + L_{max})\rho^k \|x^* - x^0\|^2}{2\epsilon}.$$

Thus we have the following high-probability bound.

Corollary 3.2. *Suppose the assumptions in Theorem 3.1 hold. Then for any $\epsilon > 0$ and $\delta \in (0, 1)$, we have*

$$\mathbf{Pr}\big(f(x^k) - f(x^\star) \leq \epsilon\big) \geq 1 - \delta$$

provided that the number of iterations k satisfies

$$k \geq \log\left(\frac{(M + L_{max})\|x^* - x^0\|^2}{2\delta\epsilon}\right) \Big/ \log\left(\frac{1}{\rho}\right).$$

Based on Theorem 3.1 and its proof, we give a deeper and stronger result that the PROXTONE achieves a linear convergence rate in the solution.

Theorem 3.3. *Suppose $\nabla g_i(x)$ and $\nabla^2 g_i$ are Lipschitz continuous with constant $L_i > 0$ and $K_i > 0$ respectively for $i = 1, ..., n$, $h(x)$ is strongly convex with $\mu_h \geq 0$. If $H^i_{\theta_{i,k}} = \nabla^2 g_i(x^{\theta_{i,k}})$ and $L_i I \preceq mI \preceq H_k^i \preceq MI$, then PROXTONE converges exponentially to $x^\star$ in expectation:*

$$\mathbb{E}[\big\|x^{k+1} - x^\star\big\|] \leq (\frac{K_{avg} + 2L_{max}}{m}\frac{M + L_{max}}{2\mu_h + m} + \frac{2L_{max}}{m})[\frac{1}{n}\frac{M + L_{max}}{2\mu_h + m} + (1 - \frac{1}{n})]^{k-1}\|x^* - x^0\|^2.$$

In order to satisfy $\mathbb{E}[\big\|x^{k+1} - x^\star\big\|] \leq \epsilon$, the number of iterations k needs to satisfy

$$k \geq (\log \rho)^{-1} \log[\frac{\epsilon}{C\|x^* - x^0\|^2}],$$

where ρ is as before and $C = \frac{K_{avg}+2L_{max}}{m}\frac{M+L_{max}}{2\mu_h+m} + \frac{2L_{max}}{m}$.

Due to the Markov's inequality, Theorem 3.3 implies the following result.

Corollary 3.4. *Suppose the assumptions in Theorem 3.3 hold. Then for any $\epsilon > 0$ and $\delta \in (0, 1)$, we have*

$$\mathbf{Pr}\big(\big\|x^{k+1} - x^\star\big\| \geq \epsilon\big) \geq 1 - \delta$$

provided that the number of iterations k satisfies

$$k \geq \log\left(\frac{((K_{avg} + 2L_{max})(M + L_{max}) + 2L_{max}(2\mu_h + m))\|x^* - x^0\|^2}{m(2\mu_h + m)\delta\epsilon}\right) \Big/ \log\left(\frac{1}{\rho}\right).$$

4 Conclusions

This paper introduces a proximal stochastic method called PROXTONE for minimizing regularized finite sums. For smooth and strongly convex problems, we show that PROXTONE not only enjoys the same linear rates as those of MISO, SAG, Prox-SVRG and Prox-SDCA, but also prove that the *solution* of this method converges in exponential rate too. There are some directions that the current study can be extended. In this paper, we have focused on the theory of PROXTONE; it would be meaningful to also do the numerical evaluation and implementation details [9]. Second, combine with randomized block coordinate method [4] for minimizing regularized convex functions with a huge number of varialbes/coordinates. Moreover, due to the trends and needs of big data, we are designing distributed/parallel PROXTONE for real life applications. In a broader context, we believe that the current paper could serve as a basis for examining the method on the proximal stochastic methods that employ second order information.

Appendix

In this Appendix, we give the proofs of the two propositions.

A Proof of Theorem 3.1

Since in each iteration of the PROXTONE, we obtain a quadratic function $g_i^k(x)$ with random parameters to approximate each $g_i(x)$:

$$g_i^k(x) = g_i(x^{\theta_{i,k}}) + (x - x^{\theta_{i,k}})^T \nabla g_i(x^{\theta_{i,k}}) + \frac{1}{2}(x - x^{\theta_{i,k}})^T H^i_{\theta_{i,k}}(x - x^{\theta_{i,k}}), \tag{A.1}$$

where $\theta_{i,k}$ is a random variable which have the following conditional probability distribution in each iteration:

$$\mathbb{P}(\theta_{i,k} = k|j) = \frac{1}{n} \quad \text{and} \quad \mathbb{P}(\theta_{i,k} = \theta_{i,k-1}|j) = 1 - \frac{1}{n}, \tag{A.2}$$

that yields

$$\mathbb{E}[\|x^* - x^{\theta_{i,k}}\|^2] = \frac{1}{n}\mathbb{E}[\|x^* - x^k\|^2] + (1 - \frac{1}{n})\mathbb{E}[\|x^* - x^{\theta_{i,k-1}}\|^2]. \tag{A.3}$$

Since $0 \preceq H^i_{\theta_{i,k}} \preceq MI$ and $\nabla^2 g_i^k(x) = H^i_{\theta_{i,k}}$, by Theorem 2.1.6 of [5] and the assumption, $\nabla g_i^k(x)$ and $\nabla g_i(x)$ are Lipschitz continuous with constant M and L_i respectively, and further $\nabla g_i^k(x) - \nabla g_i(x)$ is Lipschitz continuous with constant $M + L_i$ for $i = 1, \ldots, n$. This together with (1.7) yieds

$$|[g_i^k(x) - g_i(x)] - [g_i^k(y) - g_i(y)] - \nabla[g_i^k(y) - g_i(y)]^T(x - y)| \leq \frac{M + L_i}{2}\|x - y\|^2.$$

Applying the above inequality with $y = x^{\theta_{i,k}}$, and using the fact that $\nabla[g_i^k(x^{\theta_{i,k}})] = \nabla[g_i(x^{\theta_{i,k}})]$ and $g_i^k(x^{\theta_{i,k}}) = g_i(x^{\theta_{i,k}})$, we have

$$|g_i^k(x) - g_i(x)| \leq \frac{M + L_i}{2}\|x - x^{\theta_{i,k}}\|^2.$$

Summing over $i = 1, \ldots, n$ yields

$$[G^k(x) + h(x)] - [g(x) + h(x)] \leq \frac{1}{n}\sum_{i=1}^{n} \frac{M + L_i}{2}\|x - x^{\theta_{i,k}}\|^2. \tag{A.4}$$

Then by the Lipschitz continuity of $\nabla g_i(x)$ and the assumption $L_i I \preceq mI \preceq H_k^i$, we have

$$\begin{aligned} g_i(x) &\leq g_i(x^{\theta_{i,k}}) + \nabla g_i(x^{\theta_{i,k}})^T(x - x^{\theta_{i,k}})| + \frac{L_i}{2}\|x - x^{\theta_{i,k}}\|^2 \\ &\leq g_i(x^{\theta_{i,k}}) + (x - x^{\theta_{i,k}})^T \nabla g_i(x^{\theta_{i,k}}) + \frac{1}{2}(x - x^{\theta_{i,k}})^T H^i_{\theta_{i,k}}(x - x^{\theta_{i,k}}) = g_i^k(x), \end{aligned}$$

and thus, by summing over i yields $g(x) \leq G^k(x)$, and further by the optimality of x^{k+1}, we have

$$f(x^{k+1}) \leq G^k(x^{k+1}) + h(x^{k+1}) \leq G^k(x) + h(x) \leq f(x) + \frac{1}{n}\sum_{i=1}^{n} \frac{M + L_i}{2}\|x - x^{\theta_{i,k}}\|^2 \tag{A.5}$$

Since $mI \preceq H_{\theta_{i,k}}$ and $\nabla^2 g_i^k(x) = H_{\theta_{i,k}}$, by Theorem 2.1.11 of [5], $g_i^k(x)$ is m-strongly convex. Since $G^k(x)$ is the average of $g_i^k(x)$, thus $G^k(x) + h(x)$ is $(m + \mu_h)$-strongly convex, we have

$$\begin{aligned} f(x^{k+1}) + \frac{m + \mu_h}{2}\|x - x^{k+1}\|^2 &\leq G^k(x^{k+1}) + h(x^{k+1}) + \frac{m + \mu_h}{2}\|x - x^{k+1}\|^2 \\ &\leq G^k(x) + h(x) \\ &= f(x) + [G^k(x) + h(x) - f(x)] \\ &\leq f(x) + \frac{1}{n}\sum_{i=1}^{n} \frac{M + L_i}{2}\|x - x^{\theta_{i,k}}\|^2. \end{aligned}$$

By taking the expectation of both sides and let $x = x^*$ yields

$$\mathbb{E}[f(x^{k+1})] - f^* \leq \mathbb{E}[\frac{1}{n}\sum_{i=1}^{n} \frac{M + L_i}{2}\|x^* - x^{\theta_{i,k}}\|^2] - \mathbb{E}[\frac{m + \mu_h}{2}\|x^* - x^{k+1}\|^2].$$

We have

$$\frac{\mu_h}{2}\|x^{k+1} - x^*\|^2 \leq \mathbb{E}[f(x^{k+1})] - f^* \leq \mathbb{E}[\frac{1}{n}\sum_{i=1}^{n} \frac{M + L_{max}}{2}\|x - x^{\theta_{i,k}}\|^2] - \mathbb{E}[\frac{m + \mu_h}{2}\|x - x^{k+1}\|^2].$$

thus

$$\|x^{k+1} - x^*\|^2 \leq \frac{M + L_{max}}{2\mu_h + m}\mathbb{E}[\frac{1}{n}\sum_{i=1}^{n} \|x^* - x^{\theta_{i,k}}\|^2]. \tag{A.6}$$

then we have

$$\begin{aligned} \mathbb{E}[\frac{1}{n}\sum_{i=1}^{n} \|x^* - x^{\theta_{i,k}}\|^2] &= \frac{1}{n}\|x^k - x^*\|^2 + (1 - \frac{1}{n})\mathbb{E}[\frac{1}{n}\sum_{i=1}^{n} \|x^* - x^{\theta_{i,k-1}}\|^2] \\ &\leq \frac{1}{n}\|x^k - x^*\|^2 + (1 - \frac{1}{n})\mathbb{E}[\frac{1}{n}\sum_{i=1}^{n} \|x^* - x^{\theta_{i,k-1}}\|^2] \\ &\leq [\frac{1}{n}\frac{M + L_{max}}{2\mu_h + m} + (1 - \frac{1}{n})]\mathbb{E}[\frac{1}{n}\sum_{i=1}^{n} \|x^* - x^{\theta_{i,k-1}}\|^2] \\ &\leq [\frac{1}{n}\frac{M + L_{max}}{2\mu_h + m} + (1 - \frac{1}{n})]^k\mathbb{E}[\frac{1}{n}\sum_{i=1}^{n} \|x^* - x^{\theta_{i,0}}\|^2] \\ &\leq [\frac{1}{n}\frac{M + L_{max}}{2\mu_h + m} + (1 - \frac{1}{n})]^k\|x^* - x^0\|^2. \end{aligned}$$

Thus we have $\mathbb{E}[f(x^{k+1})] - f^* \leq \frac{M+L_{max}}{2}[\frac{1}{n}\frac{M+L_{max}}{2\mu_h+m} + (1 - \frac{1}{n})]^k\|x^* - x^0\|^2$.

B Proof of Theorem 3.3

We first examine the relations between the search directions of Prox-N and PROXTONE.

By (1.4), (1.5) and Fermat's rule, Δx_{Prox-N}^k and Δx^k are also the solutions to

$$\begin{aligned} \Delta x_{Prox-N}^k &= \arg\min_{d \in \mathbf{R}^p} \nabla g(x^k)^T d + (\Delta x_{Prox-N}^k)^T H_k d + h(x^k + d), \\ \Delta x^k &= \arg\min_{d \in \mathbf{R}^p} (\nabla_k + H_k x^k)^T d + (\Delta x^k)^T H_k d + h(x^k + d). \end{aligned}$$

Hence Δx^k and Δx^k_{Prox-N} satisfy

$$\begin{aligned}&\nabla g(x^k)^T\Delta x^k + (\Delta x^k_{Prox-N})^T H_k \Delta x^k + h(x^k+\Delta x^k)\\ \geq &\nabla g(x^k)^T\Delta x^k_{Prox-N} + (\Delta x^k_{Prox-N})^T H_k \Delta x^k_{Prox-N} + h(x^k+\Delta x^k_{Prox-N})\end{aligned}$$

and

$$\begin{aligned}&(\nabla_k + H_k x^k)^T\Delta x^k_{Prox-N} + (\Delta x^k)^T H_k \Delta x^k_{Prox-N} + h(x^k+\Delta x^k_{Prox-N})\\ \geq &(\nabla_k + H_k x^k)^T\Delta x^k + (\Delta x^k)^T H_k \Delta x^k + h(x^k+\Delta x^k).\end{aligned}$$

We sum these two inequalities and rearrange to obtain

$$\begin{aligned}&(\Delta x^k)^T H_k \Delta x^k - 2(\Delta x^k_{Prox-N})^T H_k \Delta x^k + (\Delta x^k_{Prox-N})^T H_k \Delta x^k_{Prox-N}\\ &\leq (\nabla_k + H_k x^k - \nabla g(x^k))^T(\Delta x^k_{Prox-N} - \Delta x^k).\end{aligned}$$

The assumptions $mI \preceq H_{\theta_{i,k}}$ yields that $mI \preceq H_k$, together with we have

$$m\|\Delta x^k - \Delta x^k_{Prox-N}\|^2 \leq \|\frac{1}{n}\sum_{i=1}^n(\nabla g_i(x^{\theta_{i,k}}) - \nabla g_i(x^k) - (x^{\theta_{i,k}} - x^k)^T H^i_{\theta_{i,k}})\|\|(\Delta x^k - \Delta x^k_{Prox-N})\|.$$

Since we have

$$\|\Delta x^k - \Delta x^k_{Prox-N}\| \leq \frac{K_{max}}{2mn}\sum_{i=1}^n \|x^{\theta_{i,k-1}} - x^k\|^2$$

Since the Prox-N method converges q-quadratically (*cf.* Theorem 3.3 of [2]),

$$\begin{aligned}\left\|x^{k+1} - x^\star\right\| &\leq \left\|x^k + \Delta x^k_{Prox-N} - x^\star\right\| + \left\|\Delta x^k - \Delta x^k_{Prox-N}\right\|\\ &\leq \frac{K_{avg}}{m}\left\|x^k - x^\star\right\|^2 + \left\|\Delta x^k - \Delta x^k_{Prox-N}\right\|,\end{aligned}$$

where Δx^k_{Prox-N} denotes the Prox-N search direction.

Combine by we have almost surely that

$$\begin{aligned}\left\|x^{k+1} - x^\star\right\| &\leq \frac{L_2}{m}\left\|x^k - x^\star\right\|^2 + \frac{L_{max}}{2mn}\sum_{i=1}^n \|x^{\theta_{i,k-1}} - x^k\|^2\\ &\leq \frac{K_{avg}}{m}\left\|x^k - x^\star\right\|^2 + \frac{L_{max}}{mn}\sum_{i=1}^n 2\|x^{\theta_{i,k-1}} - x^*\|^2 + \frac{L_{max}}{mn}\sum_{i=1}^n 2\|x^* - x^k\|^2.\end{aligned}$$

Then by (A.6), we have

$$\left\|x^{k+1} - x^\star\right\| \leq (\frac{K_{avg}+2L_{max}}{m}\frac{M+L_{max}}{2\mu_h+m} + \frac{2L_{max}}{m})\mathbb{E}[\frac{1}{n}\sum_{i=1}^n \|x^{\theta_{i,k}} - x^*\|^2]$$

which yieds

$$\left\|x^{k+1} - x^\star\right\| \leq (\frac{K_{avg}+2L_{max}}{m}\frac{M+L_{max}}{2\mu_h+m} + \frac{2L_{max}}{m})[\frac{1}{n}\frac{M+L_{max}}{2\mu_h+m} + (1-\frac{1}{n})]^k\|x^* - x^0\|^2.$$

References

[1] Dimitri P Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: a survey. *Optimization for Machine Learning*, 2010:1–38, 2011.

[2] Jason Lee, Yuekai Sun, and Michael Saunders. Proximal newton-type methods for convex optimization. In *Advances in Neural Information Processing Systems*, pages 836–844, 2012.

[3] Julien Mairal. Optimization with first-order surrogate functions. *arXiv preprint arXiv:1305.3120*, 2013.

[4] Yu Nesterov. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM Journal on Optimization*, 22(2):341–362, 2012.

[5] Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, Boston, 2004.

[6] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *arXiv preprint arXiv:1309.2388*, 2013.

[7] Shai Shalev-Shwartz and Tong Zhang. Proximal stochastic dual coordinate ascent. *arXiv preprint arXiv:1211.2717*, 2012.

[8] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss. *The Journal of Machine Learning Research*, 14(1):567–599, 2013.

[9] Jascha Sohl-Dickstein, Ben Poole, and Surya Ganguli. Fast large-scale optimization by unifying stochastic gradient and quasi-newton methods. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 604–612, 2014.

[10] Lin Xiao and Tong Zhang. A proximal stochastic gradient method with progressive variance reduction. *arXiv preprint arXiv:1403.4699*, 2014.